
# Towards realizing disaggregated datacenters of the future

Over the last years, the number and size of Cloud infrastructures have experienced a rapid increment. Typically, public Cloud services are hosted inside datacenter (DC) infrastructures, which, in turn, house all the required hardware resources to run these services. Traditional datacenters have a relatively static computing architecture, consisting of a number of servers, each with a fixed number of CPUs and RAM and with potentially different types of hardware accelerators (Figure (a)). Datacenter operators have used this monolithic server model for years, however, as the variety of hosted applications, the hardware heterogeneity and the adoption of cloud computing increase, so does the complexity of operating efficiently such an infrastructure. Typical reasons, among others, are the limited resource utilization of modern infrastructures, the difficulty of integrating new HW devices and the handling of HW failures.

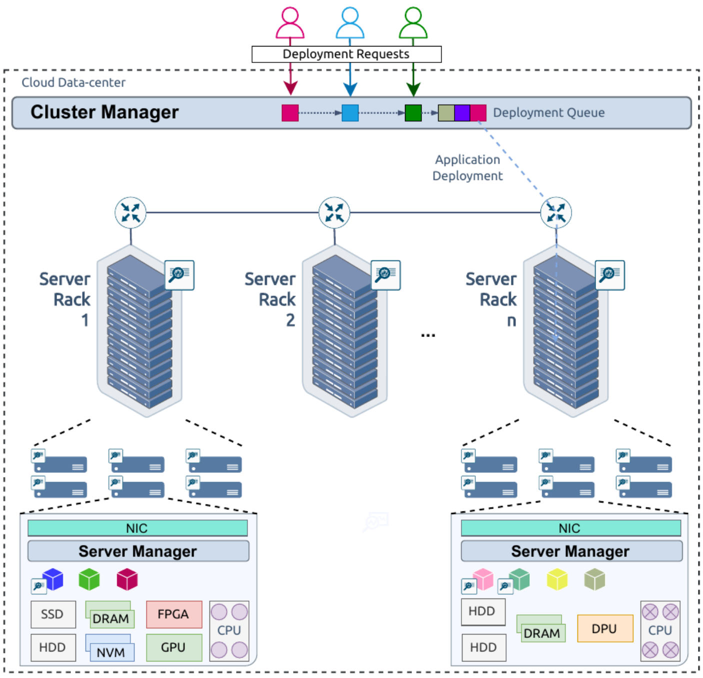


Figure (a): Traditional DC architecture

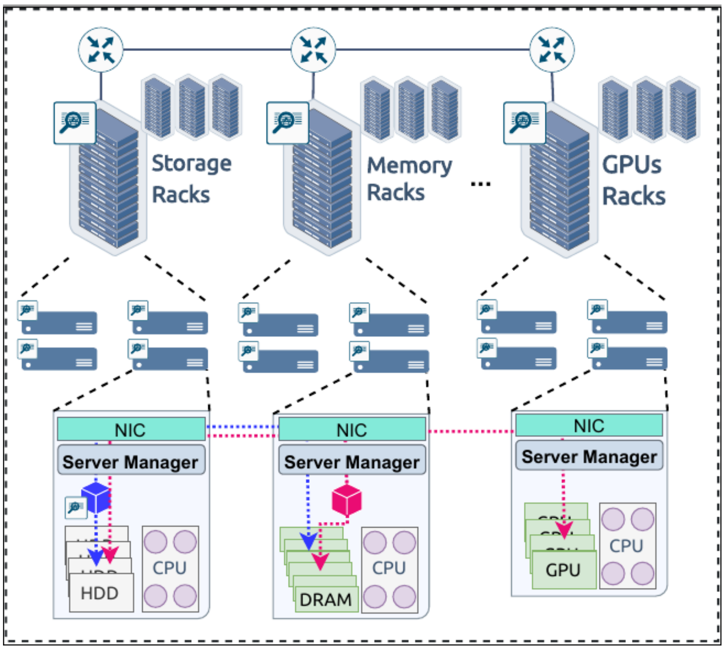


Figure (b): Disaggregated DC architecture

Towards realizing infrastructures of scale, datacenter operators are moving towards a new computing paradigm, referenced in the literature as *hardware resource disaggregation* or *composable disaggregated infrastructures*. Figure (b) shows a simplified overview of the hardware disaggregation concept. As shown, in a hardware disaggregated system, there is a transformation of general-purpose monolithic servers into network-attached resource pools that may be constructed, managed, and scaled separately. Similar to conventional datacenters, servers are organized inside racks, which, however, are resource-specific (e.g., storage, memory, etc.). Each server on the rack combines a pile of a particular set of resources with CPUs to manage them and an interface to a specialized network fabric to communicate with other resources on the cluster.

In this direction, a set of diploma theses will be examined targeting the novel concept of hardware disaggregation. Depending on the background and interests of the student, the thesis can expand towards:

1. Exploring AI-driven resource management schemes for efficiently selecting appropriate set of computing resources (e.g., local/remote memory) for deployed applications.
2. Building simulators that provide realistic imitation of the controls and operation of hardware disaggregated systems.
3. Exploring the impact of memory disaggregation on state-of-the-art in-memory applications (e.g., Redis, Spark) and extending the frameworks to efficiently leverage local/remote memory tiers.

## USEFUL PREREQUISITES

- Linux, Bash/Shell scripting.
- Python, Machine Learning frameworks and libraries

## USEFUL LINKS

- https://research.ibm.com/projects/composable-disaggregated-infrastructure
- https://engineering.purdue.edu/WukLab/disaggregate.html

## RELATED MATERIAL


1. Shan, Yizhou, et al. "{LegoOS}: A Disseminated, Distributed {OS} for Hardware Resource Disaggregation." 13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18). 2018.
2. C. Pinto et al., "ThymesisFlow: A Software-Defined, HW/SW co-Designed Interconnect Stack for Rack-Scale Memory Disaggregation," 2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), Athens, Greece, 2020, pp. 868-880, doi: 10.1109/MICRO50266.2020.00075.
3. Ruan, Zhenyuan, et al. "{AIFM}:{High-Performance},{Application-Integrated} Far Memory." 14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20). 2020.
4. Kumar, Gautam, et al. "Swift: Delay is simple and effective for congestion control in the datacenter." Proceedings of the Annual conference of the ACM Special Interest Group on Data Communication on the applications, technologies, architectures, and protocols for computer communication. 2020.
5. P. Koutsovasilis, M. Gazzetti and C. Pinto, "A Holistic System Software Integration of Disaggregated Memory for Next-Generation Cloud Infrastructures," 2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid), Melbourne, Australia, 2021, pp. 576-585, doi: 10.1109/CCGrid51090.2021.00067.


## CONTACT INFORMATION


- Dimosthenis Masouros, PhD candidate Microlab NTUA (dmasouros@microlab.ntua.gr)
- Christian Pinto (IBM Research Europe - Ireland)
- Sotirios Xydis, Assistant Professor Microlab NTUA (sxydis@microlab.ntua.gr)
- Dimitrios Soudris, Professor Microlab NTUA (dsoudris@microlab.ntua.gr)